

# Short Portfolio Training for CSP Solving

Mirko Stojadinović · Mladen Nikolić · Filip Marić

Received: date / Accepted: date

**Abstract** Many different approaches for solving Constraint Satisfaction Problems (CSPs) and related Constraint Optimization Problems (COPs) exist. However, there is no single solver (nor approach) that performs well on all classes of problems and many portfolio approaches for selecting a suitable solver based on simple syntactic features of the input CSP instance have been developed. In this paper we first present a simple portfolio method for CSP based on k-nearest neighbors method. Then, we propose a new way of using portfolio systems — training them shortly in the exploitation time, specifically for the set of instances to be solved and using them on that set. Thorough evaluation has been performed and has shown that the approach yields good results. We evaluated several machine learning techniques for our portfolio. Due to its simplicity and efficiency, the selected k-nearest neighbors method is especially suited for our short training approach and it also yields the best results among the tested methods. We also confirm that our approach yields good results on SAT domain.

**Keywords** CSP Portfolio · Short Training · k-nearest neighbors · SAT Portfolio

## 1 Introduction

*Constraint satisfaction problems (CSPs)* and related *Constraint optimization problems (COPs)* (Apt 2003; Rossi et al. 2006) over finite domains are wide classes of problems that include many problems relevant for real world applications (e.g., scheduling, timetabling, sequencing, routing, rostering, planning) (Rossi et al. 2006). Many different approaches for solving

---

Mirko Stojadinović (✉)  
Faculty of Mathematics, University of Belgrade, Serbia  
E-mail: mirkos@matf.bg.ac.rs

Mladen Nikolić  
Faculty of Mathematics, University of Belgrade, Serbia  
E-mail: nikolic@matf.bg.ac.rs

Filip Marić  
Faculty of Mathematics, University of Belgrade, Serbia  
E-mail: filip@matf.bg.ac.rs

CSPs exist (e.g., constraint propagation, backtracking search algorithms, local search methods, constraint logic programming, operation research methods, answer set programming, reduction to SAT/SMT, lazy clause generation) (Rossi et al. 2006) and there are many state-of-the-art solvers that implement these approaches.

It is well-known that there is neither single solver nor single approach suitable for all problems. When solving a CSP instance, one should consider using several solvers and several different configurations (setups) of each solver (if the solver is configurable). If a multiprocessor machine is available, different solvers could be run in parallel, until one of them solves the problem. However, in many cases this is not feasible or desirable, so it is preferable to somehow guess the solver that would give the best results for each given instance and then run only that solver. *Portfolio approaches* that have been successfully used for SAT (e.g., (Xu et al. 2008; Nikolić et al. 2009, 2011; Malitsky et al. 2011; Kadioglu et al. 2011; Malitsky and Sellmann 2012; Malitsky et al. 2012, 2013)) but also for CSP (e.g., (O'Mahony et al. 2008; Kiziltan et al. 2011; Amadini et al. 2013, 2014; Hurley et al. 2014)) assume that a number of different solvers are available and for each input instance these approaches select a solver (and its configuration) that should be run. This choice usually consists of applying some machine learning technique and is most often based on the knowledge gained during previous runs of the available solvers on some training instances and on some syntactic characteristics of the instance to be solved.

*Finite linear CSP* (Tamura and Banbara 2008) is a special class of constraint satisfaction problems that is often encountered in applications. We consider only that class of problems, but also allow global constraints (Beldiceanu et al. 2005). In our previous work (Stojadinović and Marić 2014) we adapted the ArgoSmArT-kNN portfolio (Nikolić et al. 2011), originally developed for SAT, to constraint satisfaction problems. The main idea behind portfolio retains the simplicity of the original approach, but the portfolio uses features of CSP instances and selects CSP solver for the input CSP instance. In that paper, we considered only CSP solvers based on reduction to SAT and we applied the portfolio approach for selecting between different *SAT encodings* that can be used for that reduction. In this paper, we select between a very wide range of available solvers, including solvers based on *reduction to SAT* (Biere et al. 2009), *reduction to SMT* (Biere et al. 2009), *lazy clause generation solvers* (Ohrimenko et al. 2009), and *constraint propagation solvers* (Rossi et al. 2006). As different solving methods are usually good at solving different types of problems, we want to exploit this fact, thus increasing the potential efficiency of portfolio. We compare the efficiency of our approach to the other state-of-the-art CSP portfolios.

One of the problems with off-the-shelf application of portfolios is that the portfolio may be trained on a set of instances with properties significantly different than the set of instances to be solved. In such cases the portfolio may perform poorly. Also, retraining the portfolio for some specific set of instances may require large amounts of time. We propose a new way of using portfolios which relies on short portfolio training. Consider a scenario in which a user wants to solve a specific, fixed set of instances as fast as possible using solvers at his disposal. In our approach, a short training run is performed on all of the given instances in order to train a portfolio, and then, the portfolio is used to solve those very same instances. Usually, when considering portfolios (or machine learning methods and their applications in general), one is concerned with their ability to generalize and would not evaluate a portfolio on the instances it was trained on, as it would give too optimistic estimate of its future performance. However, if the goal is just to solve a specific set of instances, one need not bother with portfolio's generalization ability and can just apply it to the instances to be solved. To our knowledge, the portfolios have not been used in this way before, mainly due

to significant times required by the standard portfolio training. Therefore, we provide a new way of applying portfolios to solve practical problems.

We experimentally evaluate effectiveness of different machine learning techniques when using short training in CSP portfolios. We evaluate portfolios based on the k-nearest neighbors method, support-vector machines, and linear regression.

Finally, we want to test if the main conclusions of our work transfer to the SAT case.

Contributions of this work are the following.

- We present a thorough experimental evaluation of CSP solvers obtained by applying 15 very diverse CSP solvers on the corpus containing more than 8,000 CSP instances.
- We formulate  $k$ -NN-based CSP portfolio and compare its effectiveness to other state-of-the-art portfolios – our portfolio gives comparable results.
- We test the effect of the training time on our  $k$ -NN-based portfolio and show that it significantly improves performance over the best-fixed solver, even when extremely low solving timeouts are used (only several seconds per instance).
- We propose a new way of using portfolio methods which relies on performing short training on the very same instances that should be solved.
- We test effectiveness of using different machine learning techniques with short training and demonstrate that  $k$ -NN-based portfolio gives the best results.
- We show that main conclusions of our work on CSP transfer also to the case of SAT portfolios.

*Overview of the paper.* In Section 2 we give some basic definitions, describe different solving methods and solvers, and present some of the most well-known portfolio approaches. In Section 3 we present experimental evaluation of different CSP solvers on a large corpus of CSP instances. The results in Section 3 are used as a basis for the next 3 sections. In Section 4 we describe our approach, compare it with other most successful approaches and estimate the effect of the training time on the portfolio effectiveness. The problem of obtaining the best results on a fixed set of instances is addressed in Section 5. This section also provides comparison of efficiency of different machine learning techniques using short training and evaluation of the efficiency of our approach on SAT problems. In Section 6 we draw some final conclusions and present ideas for further work.

## 2 Background

In this section we introduce key notions used in the rest of the paper and we analyze prior results in this area.

### 2.1 Finite Linear CSP

**Definition 1** *Linear expressions* over the set of integer variables  $V$  are algebraic expressions of the form  $\sum_{k=1}^n a_k x_k$  where all  $x_k$  are variables from  $V$  and all  $a_k$  are integers.

A *Finite Linear CSP in CNF* is a tuple  $(V, L, U, B, S)$  where

1.  $V$  is a finite set of integer variables,
2. functions  $L : V \mapsto \mathbb{Z}$  and  $U : V \mapsto \mathbb{Z}$  give lower and upper bound of integer variables and these bounds determine the domain  $D(x)$  of each variable  $x$ ,
3.  $B$  is a set of Boolean variables,

4.  $S$  is a finite set of clauses (over  $V$  and  $B$ ). Clauses are formed as disjunctions of literals where literals are the elements of the union of the sets  $B$ ,  $\{\neg p \mid p \in B\}$  and  $\{e \leq c \mid e \text{ is linear expression over } V, c \in \mathbb{Z}\}$ .

A *Solution* of Finite Linear CSP in CNF is an assignment of Boolean values to Boolean variables and integer values to integer variables satisfying their domains, such that all clauses from  $S$  are satisfied when variables are replaced by their values.

*Example 1* A solution of Finite Linear CSP  $V = \{x_1, x_2, x_3\}$ ,  $L = \{x_1 \mapsto 1, x_2 \mapsto 1, x_3 \mapsto 2\}$ ,  $U = \{x_1 \mapsto 2, x_2 \mapsto 4, x_3 \mapsto 3\}$ ,  $B = \{p\}$ ,  $C = \{p \vee x_1 + x_3 \leq 4, \neg p \vee x_3 + (-1) \cdot x_1 \leq 0, x_1 \leq 1 \vee 2 \cdot x_2 \leq 4\}$  is the assignment  $\{p \mapsto \perp, x_1 \mapsto 1, x_2 \mapsto 3, x_3 \mapsto 2\}$ .

In applications, the input syntax is usually modified so that it allows non-contiguous domains, formulae with arbitrary Boolean structure (not only CNF) and with literals formed by applying other arithmetic operations (e.g., integer division, modulo) and other arithmetic relations (e.g.,  $<$ ,  $\geq$ ,  $>$ ,  $=$ ). All these formulae alongside the clauses described in Definition 1 are called *intensional constraints*. Another usual modification of the syntax is the usage of *extensional constraints* (sometimes called *user-defined relations*) that are defined by a table of allowed/disallowed assignments to the variables that they constrain. A *global constraint*<sup>1</sup> is a constraint that encapsulates a set of other constraints with two main purposes: to increase expressiveness and improve efficiency by allowing specialized algorithms for processing these constraints (Régin 2004). Intensional, extensional, and global constraints can be reduced to finite linear CSP in CNF form during preprocessing, but usually more efficient procedures are obtained if these are treated directly.

*Example 2* Constraint solver Sugar (Tamura and Banbara 2008) solves finite linear CSP by reduction to SAT and it uses very simple input language. We give here an example of finite linear CSP specification in this language.

```
(int x1 1 2) (int x2 1 4) (int x3 2 3)
(imp (>= (+ x1 (* 2 x3)) 3) (and (< x1 x2) (<= x3 (+ x1 x2))))
(alldifferent x1 x2 x3)
```

The example uses intensional constraints and global constraint `alldifferent`, stating that all its arguments have to take mutually different values. The first row declares the domains of the variables and the rest impose constraints on these variables. One of the solutions to this problem is the assignment  $x_1 = 1, x_2 = 2, x_3 = 3$ .

## 2.2 Modeling languages

Before solving, a constraint satisfaction problem must be somehow specified. For this purpose, many modeling languages exist. Two most common modeling languages are *MiniZinc* and *XCSP*. We also consider the *Sugar* language used by several solvers in our portfolio.

MiniZinc (Nethercote et al. 2007) is a high-level constraint modeling language. Before solving, most solvers compile MiniZinc specifications to a low-level target language FlatZinc. G12 MiniZinc distribution contains many tools and solvers accepting MiniZinc and FlatZinc formats.

<sup>1</sup> A catalogue of global constraints (Beldiceanu et al. 2005) is available online: <http://www.emn.fr/z-info/sdemasse/gccat>

XCSP (Roussel and Lecoutre 2009) is an XML-like low-level format used in several CSP solving competitions.

Instances from XCSP format can be directly translated to Sugar input format<sup>2</sup>, that has much simpler syntax (Example 2).

## 2.3 Solving methods for CSPs

*Propagation-based systems.* The process of *Constraint propagation* (Rossi et al. 2006) reduces a CSP to an equivalent problem, simpler to solve. The reduction is most often done by removing values from the domains of the variables that cannot be the part of any solution to the problem. Many solvers that are based on constraint propagation techniques were developed. These solvers use search heuristics and algorithms for performing constraint propagation, thus achieving different types of consistencies (Apt 2003; Rossi et al. 2006). Solvers *Mistral* (Hebrard 2008) and *Abscon* (Merchez et al. 2001) that participated in CSP competitions belong to this type of solvers.

*Reduction to SAT.* *Propositional satisfiability problem (SAT)* (Biere et al. 2009) is the problem of deciding if there is a truth assignment under which a given propositional formula (in conjunctive normal form) evaluates to true. It is a canonical NP-complete problem (Cook 1971) and it holds a central position in the field of computational complexity. When using reduction to SAT, CSP instances are encoded as SAT instances and modern efficient satisfiability solvers are used for finding solutions that are then converted back to the solutions of the original CSPs.

A fundamental design choice when encoding finite domain constraints into SAT concerns the representation of integer variables. Several different encoding schemes (Tanjo et al. 2011; Stojadinović and Marić 2014) have been proposed and successfully used in various applications. There are several tools that reduce CSPs to SAT using one or more of standard encodings (e.g., direct encoding, support encoding, order encoding, log encoding).

Sugar is a constraint solver that solves finite linear CSPs by translating them into SAT using the order encoding method (Tamura et al. 2009) and then solving SAT instances by several supported SAT solvers.

Azucar (Tanjo et al. 2012) is a successor of Sugar that uses the compact-order encoding (Tanjo et al. 2011) for translating finite linear CSP into SAT. It is tuned for solving specific large domain sized CSP instances. Log encoding is a special case of compact-order encoding so Azucar can also use this encoding when reducing to SAT.

*meSAT* (Stojadinović and Marić 2014) (Multiple Encodings of CSP to SAT) is a system using different encodings and their combinations. The supported encodings are: direct (Walsh 2000), support (Gent 2002), direct-support (Stojadinović and Marić 2014), order (Tamura et al. 2009) and direct-order (Stojadinović and Marić 2014).

*Reduction to SMT.* *Satisfiability modulo theories (SMT)* (Biere et al. 2009) is a research field concerned with the satisfiability of formulae with respect to some decidable first-order background theory (or combination of them). Some of these theories are *Linear Integer Arithmetic*, *Integer Difference Logic*, *Linear Real Arithmetic*, etc. There are several systems that solve CSPs by reduction to SMT (e.g. *fzn2smt* (Bofill et al. 2010)).

---

<sup>2</sup> Web page <http://bach.istc.kobe-u.ac.jp/sugar/current/docs/syntax.html> contains full syntax of Sugar input language.

*Lazy clause generation.* In the *lazy clause generation* approach (Ohrimenko et al. 2009), finite domain constraint propagation engine is combined with a SAT solver: propagators are mapped into clauses and passed to SAT solver, which uses unit propagation and then returns obtained information back to the engine. Contrary to the eager approach, clauses are not generated a priori but are constructed and given to the SAT solver during the solving phase. The lazy propagation approach can be viewed as a special form of Satisfiability Modulo Theories solver, where each propagator is considered as a separate theory, and theory propagation is used to learn clauses. Solvers `mzn-g12cpx` and `mzn-g12lazy` (included in the MiniZinc G12 distribution) implement lazy clause generation.

## 2.4 Machine learning techniques used for CSP portfolio

Machine learning methods are often used to predict value of some outcome variable based on the values of some other related variables, called *features*. There are various methods that address this task. Usually, some form of statistical model that expresses the dependence of the outcome variable on feature variables is assumed. The coefficients of the model are determined using a set of instances for which both the values of outcome and feature variables are known. We describe some frequently used methods that are used in this paper.

*Linear regression* is one of the most often used machine learning methods. The variable to be predicted is modeled by a linear combination of feature variables (Bishop et al. 2006). The coefficients of the model can be determined by the well known least squares method in order to minimize sum of square differences between the predicted and actual outcomes on training data.

In *k-nearest neighbors method*, it is assumed that the outcome variable is modeled based on its values on  $k$  instances from the training set that are closest to the input instance, with respect to some distance measure defined over feature vectors. Number  $k$  is selected empirically. If the outcome variable is continuous, its value is taken to be a linear combination of its values on the neighbor instances. For the coefficients of the linear combination the simplest choice is  $1/k$ , but they can also depend on the distances involved. If the outcome variable is discrete, its value is decided by voting of neighbor instances.

In its basic form, *support vector machine (SVM) regression* models the outcome variable by a linear combination of feature variables with additional requirements. Firstly, the predictions need to deviate from actual values of the outcome variable by at most  $\epsilon$  (which is a metaparameter) and secondly, the model should be as flat as possible. Flatness is ensured by minimizing the norm of the coefficient vector. Modifications are made to tackle the cases in which there is no function which fits the deviation constraints and to allow for nonlinear regression functions.

Training algorithms usually have some metaparameters which influence the model that the algorithm yields for specific data set. Their selection is a nontrivial process and the best choice of the values for metaparameters is usually made by evaluating the models that different values of metaparameters yield.

In order to check if models produced by machine learning methods generalize well, they should be evaluated on a data set different than the training set. The evaluation is usually performed in one of two ways. First one is to split the available data to training and test set, perform training and produce a model on training set, and then measure how well it predicts the outcome variable in the test set. In this scenario, the question arises how to select the test set. Different splits to training and test set may result in different evaluation results. To avoid that, second evaluation scenario —  $k$ -fold cross-validation — is used. The available

data is split to  $k$  parts (called folds). For each part, the model is trained on  $k - 1$  other parts and evaluated on that part. At the end, evaluation results are aggregated. Cross-validation produces more reliable estimate of the generalization capability of the model, but is clearly more time consuming.

## 2.5 Solver selection for SAT and CSP

*SAT portfolios.* The instance-based solver selection problem has been widely studied in the SAT community. Based on the characteristics of the input instance, either some parameters of a single solver are tuned, or one of several available solvers is selected to be applied on that instance (so called solver portfolio). The most successful results are based on machine-learning techniques (e.g., SATZilla (Xu et al. 2008), ArgoSmArT-kNN (Nikolić et al. 2009, 2011), ISAC (Malitsky and Sellmann 2012), Non-Model-Based Algorithm Portfolios for SAT (Malitsky et al. 2011)). Each SAT instance is characterized by a set of its features (most of them are purely syntactic and extracted from the CNF representation). Usually, a training corpus is solved by different SAT solvers (or a single solver configured by different parameters) and for each instance, its solving times for each solver are assigned to its feature vector. For each solver, a predictive model is learned, which describes the dependence between feature vectors and solving times. When a new instance is to be solved, the most suitable solver is chosen, based on solving times predicted by models of each solver for feature vector of the input instance.

*CSP portfolios.* Algorithm portfolios have recently been applied to constraint satisfaction problem solving. CPHYDRA (O'Mahony et al. 2008) is an algorithm portfolio for CSP that uses  $k$ -nearest neighbors algorithm to determine one or more solvers to be run on an unseen problem instance. The superiority of the portfolio over each of its constituent solvers is demonstrated using challenging benchmark problem instances from CSP Solver Competition. Another approach, described by Kiziltan et al. (Kiziltan et al. 2011), uses run-time classifiers (categories are: “short”, “medium” and “long”) to minimize the average solving time of each instance. This portfolio uses features of CPHYDRA and SATZilla and the combination of the two. Work by Amadini et al. (Amadini et al. 2013) compares efficiency of different portfolio approaches based on SAT portfolio techniques and machine learning algorithms.

A recent system SUNNY (Amadini et al. 2014) outperforms CPHYDRA and some portfolios originally developed for SAT, but adapted to CSP solving (e.g. SATZilla). For some new instance to be solved, SUNNY uses  $k$ -NN algorithm to select one or more solvers that should be applied. In practice, the number of selected solvers is very rarely greater than two. Selected solvers share available time (each selected solver is assigned certain amount of time from the total time available, based on the efficiency of each solver on  $k$  nearest instances). SUNNY uses 155 features, gathered from instances in the MiniZinc input format, using tool *mzn2feat* developed by the same authors. From these features, 11 are dynamic, i.e. obtained by running Gecode solver for 2 seconds. The original experimental evaluation (Amadini et al. 2014) included instances from CSP Solver Competition and from MiniZinc distribution, and 11 solvers from MiniZinc Challenge.

Proteus (Hurley et al. 2014) is a hierarchical portfolio that outperforms many other machine learning approaches adapted for CSP solving. When a new instance is to be solved, a solver is selected by making decisions on two or three levels. On the first level, only the solving approach is chosen – either (not SAT-based) CSP solving or reduction to SAT. If



CSP solving is chosen, the CSP solver that should be applied is selected on the second level. If reduction to SAT is chosen, only the encoding method is chosen on the second level, and the solver is selected separately, on the third level. As stated by the authors, the advantage of this approach is that most suitable technique can be used for each decision level, and this can lead to improved performance. On the other hand, usage of different techniques at each level means the system is of a much greater complexity than some other approaches. *Proteus* uses 36 CSP features, the same as CPHYDRA, and among these features some dynamic, obtained by running *Mistral* solver for 2 seconds. For each of 3 SAT encodings, *Proteus* uses 54 SAT features. Instances from CSP Solver Competition were used and *Proteus* was tested with 4 CSP solvers, 3 SAT encodings and 6 SAT solvers.

In our previous work (Stojadinović and Marić 2014), we described portfolio for selecting between different SAT encodings used by system *meSAT*. The selection process is based on  $k$ -NN algorithm and the constructed portfolio outperformed all the constituent encodings (direct-support, order, direct-order) used in the experiments. One of the assumptions of that paper was that the most efficient encoding in solving easier instances of some family of instances (coming from a single problem) will also be the most efficient in solving harder instances of the same family. On the instances used in the experiments, this assumption was demonstrated to be right and the experiments showed that portfolio *trained only on easy instances* (solvable by all used encodings in some small time) can outperform any constituent encoding.

This paper is a continuation of our previous work: the same portfolio is used, but now with different CSP solvers and not only different encodings. We provide the pseudocode of the portfolio and describe it more formally. The ways the portfolios are used in the original paper and in this paper significantly differ – in this paper there is no assumption that instances are divided in families, the training is not performed only on easy instances, the aims of experiments are different, etc.

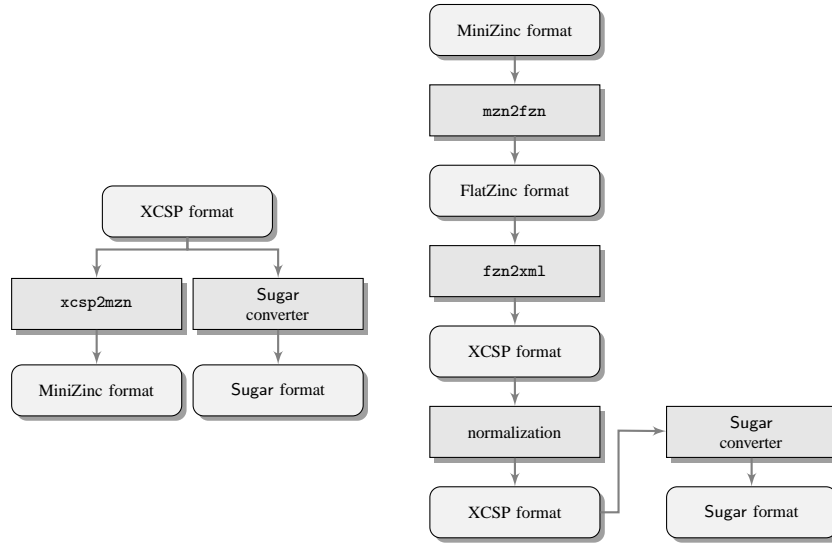
### 3 Experimental Setup

In this section, we describe our experimental setup that will be used throughout the paper. We use a rich set of available CSP solvers and a very wide corpus of available CSP benchmarks. Aside of describing them, here we also present the evaluation of these solvers on these benchmarks.

*Solvers.* For reduction to SAT, we used direct-support, order and direct-order encodings implemented in system *meSAT* 1.1 (Stojadinović and Marić 2014), order encoding implemented in system *Sugar* 2.1.3 (Tamura and Banbara 2008), and log and compact-order encoding implemented in system *Azucar* 0.2.4 (Tanjo et al. 2012). SAT solver *Minisat* 2.2 (Eén and Sörensson 2003) was used in all cases when reduction to SAT was performed. We extended *meSAT* to enable reduction of CSP instances to SMT-LIB language<sup>3</sup>. The translation of most constraints is straightforward, and only global constraints are decomposed to more simpler constraints. SMT solvers *Yices* 2.2.0 (Dutertre and de Moura 2006) and *Z3* 4.2 (de Moura and Bjørner 2008) were used for solving generated SMT-LIB instances. We used two solvers implementing constraint propagation techniques: *Abscon* 112V4 (Merchez et al. 2001) and *Mistral* 1.545 (Hebrard 2008). We also used solvers from G12 *MiniZinc* 1.6

<sup>3</sup> The source code of our implementation and the instances used in experiments (but without third-party solvers, due to specific licensing) are available online from *meSAT* web page: <http://jason.matf.bg.ac.rs/~mirkos/Mesat.html>





**Fig. 1** Conversions between different input formats.

(Nethercote et al. 2007) distribution: `mzn-g12lazy` and `mzn-g12cpx` implementing lazy clause generation, and `mzn-g12fd` and `mzn-g12mip`. Solver Gecode 4.2.1 (Schulte et al. 2006) was also used.

*Instances.* We used two publicly-available corpora of CSP instances: (i) CPAI09 corpus containing all instances used in Fourth International CSP Solver Competition<sup>4</sup>, (ii) instances from MiniZinc corpus available on MiniZinc page<sup>5</sup>. We used three formats of input files: MiniZinc language (Nethercote et al. 2007), XCSP format (Roussel and Lecoutre 2009) and Sugar (Tamura and Banbara 2008) input format.

Instances are converted between different formats using of-the-shelf tools and the conversions are presented in Figure 1. Instances from the first corpus were automatically converted from the original input language to MiniZinc by the converter `xcsp2mzn` available on MiniZinc page and to Sugar input format by the converter included in the Sugar distribution. Instances from the second corpus use model files and data files from MiniZinc distribution corpora. These instances were translated to FlatZinc format by `mzn2fzn`, then from this format to XCSP format by `fzn2xml` (both converters are available in MiniZinc distribution) and then normalized by the tool provided by Amadini et al. (Amadini et al. 2013). From XCSP format these instances were also converted to Sugar input format.

We excluded instances that could not be converted from any of two formats to any other format and instances for which there was a solver giving wrong satisfiability answer (surprisingly, there were 541 instances where at least one of the solvers gave wrong answer, and 8 out of 15 solvers gave wrong answers on some instances). The final corpus consisted of 8436 instances.

*Experimental environment.* All tests were performed on a multiprocessor machine with AMD Opteron(tm) CPU 6168 on 1.9Ghz with 2GB of RAM per CPU, running Linux. Solv-

<sup>4</sup> <http://www.cril.univ-artois.fr/CPAI09>

<sup>5</sup> <http://www.minizinc.org>

Solving method	Solver	# (out of 8436)	Time (days)
Propagation based	<i>Abscon</i>	4402	30.2
	<i>Mistral</i>	6216	16.9
Reduction to SAT	<i>meSAT</i> (direct-support)	3128	38.6
	<i>meSAT</i> (order)	3985	32.6
	<i>meSAT</i> (direct-order)	3984	32.6
	<i>Sugar</i> (order)	5548	22.7
	<i>Azucar</i> (compact-order)	4402	31.1
	<i>Azucar</i> (log)	3804	30.7
Reduction to SMT	<i>Z3</i>	4938	26.7
	<i>Yices</i>	4978	26.0
Lazy clause generation	<i>mzn-g12cpx</i>	2883	39.4
	<i>mzn-g12lazy</i>	2865	41.3
Other	<i>mzn-g12mip</i>	1544	49.0
	<i>mzn-g12fd</i>	2716	41.7
	<i>Gecode</i>	2627	41.4
	<i>best-fixed</i>	6216	16.9
	<i>oracle</i>	7631	7.0

**Table 1** Results of experimental evaluation of CSP solvers; the encoding used for reduction to SAT is given in parentheses; # denotes the number of solved instances; *best-fixed* – single solver achieving the best results (*Mistral* in this case), *oracle* – the best solver is chosen for each instance.

ing timeout was 600 seconds for each instance (for total time including selecting the solver where needed, encoding where needed, and solving).

*Solver comparison.* We exhaustively ran all solvers on all instances. The results are shown in Table 1. The number of solved instances and total solving time (in days) are shown (for each unsolved instance, the timeout value of 600 seconds is assumed). Mean time per instance is directly computable from the total time and the number of instances. We give results for two additional methods: (i) *oracle* (or virtual best) method that would select the best solver for each instance (this method is not feasible in practice since for each instance it must always guess the best solver) and (ii) the *best-fixed* method – one solver that gives the best overall performance. The results indicate that the difference between the *best-fixed* and the *oracle* is 1415 instances, so, it makes sense to apply a portfolio approach.

## 4 CSP Portfolio

In this section we present a simple portfolio for CSP. The basic principle was taken from SAT portfolio ArgoSmArT-kNN (Nikolić et al. 2011), but the new portfolio was constructed based on features (described in Section 4.1) and solvers suited for CSPs. In Section 4.2 we describe the portfolio, in Section 4.3 we compare efficiency of ArgoCSP-kNN to different, already existing, portfolios. Finally, in Section 4.4 we briefly evaluate the effect of instance solving time on portfolio effectiveness.

### 4.1 Portfolio features

Unlike some other approaches (e.g., (Kiziltan et al. 2011)) that use features of the generated SAT instances (in case of reduction to SAT), we use only features extracted from the orig-

inal CSP instance. We considered 70 different features<sup>6</sup> divided in several groups: features related to the sizes of the domains of integer variables for all variables in the instance (e.g., average domain size), and for the variables included in each different type of constraint, features related to the number of all variables and variables with non-contiguous domains, features related to the number and the percentage of the constraints of different types — intensional (e.g., percentage of intensional constraints among all the constraints), extensional, global (e.g., average arity of global constraints), as well as for each specific type of constraint (e.g., number of arithmetic constraints, number of multiplications, sum of domains of variables involved in multiplications, number of *all-different* constraints), etc.

In most of the experiments – except when we explicitly state that we did differently – we used all 70 features. The time used for the feature extraction is small (about 0.05 seconds in average on all instances).

## 4.2 ArgoCSP-kNN portfolio

Before it can be applied, ArgoCSP-kNN portfolio must be trained. *Portfolio training* consists of solving the training corpus (*the preparation*) with all available solvers and *predictive model training*. Once trained, portfolio can be used many times. That phase we call *exploitation* phase.

*Preparation.* First, the features of all training instances and performance of all solvers on these instances need to be gathered. For solving each training instance, each solver is given certain amount of time, called the *solving timeout*. Solver performance on an instance is expressed by a *PAR10* score (Hutter et al. 2009) — the solving time if the instance is solved within the given timeout, or the timeout value multiplied by 10, otherwise. Results of preparation phase are expressed by two tables. For each instance, the first table contains its features. The second table contains the *PAR10* score for the given solving timeout for each solver and each instance.

*Predictive model training.* The purpose of this phase is to select (based on the results of the preparation phase) the optimal values for  $k$  (the number of neighbors) and  $d$  (the distance measure) that are later going to be used in the portfolio exploitation. This is done by using the *cross-validation* technique. Different values of  $k$  and  $d$  are tried. For each fixed combination of these two metaparameters, a 5-fold cross-validation on prepared data is performed. At each turn, for each instance of one fold, its  $k$  nearest neighbors (with respect to the distance measure  $d$ ) are found among instances of four other folds. The solver having the smallest sum of *PAR10* scores on these neighboring instances is considered the best and selected. The combination of  $k$  and  $d$  giving the best results on all training instances (the smallest sum of *PAR10* scores of selected solvers for all training instances) is the one that is going to be used by the portfolio.

*The portfolio and its exploitation.* The solver selection algorithm (based on (Nikolić et al. 2011)) is shown in Figure 2. The values for the input parameter `preparation_data` are collected during preparation, while the values for input parameters `k` and `d` are chosen in the process of predictive model training. In the exploitation phase, first the features of the input instance are computed. Then, the nearest neighbors of the input instance in the whole

---

<sup>6</sup> The already mentioned *meSAT* web page contains a detailed description of all 70 features.

```

instance - instance to be solved
solvers - set of solvers
preparation_data - features and PAR10 scores of solvers on training instances
k - number of neighbor instances to be considered
d - distance measure

```

```

ArgoCSP-kNN (instance, solvers, preparation_data, k, d)
    features = CSPFeatures (instance);
    neighbors = nearestNeighbors (preparation_data, k, d, features);
    best_score = ∞;
    for s in solvers
        if (par10 (s, neighbors) < best_score)
            best_score = par10 (s, neighbors);
            best_solver = s;
    return best_solver;

```

**Fig. 2** ArgoCSP-kNN solver selection procedure.

training corpus are found and the best solver on these instances is used to solve the input instance. Function `par10` returns the sum of PAR10 scores of some specific solver on neighboring instances.

### 4.3 Comparison with other CSP portfolios

In this subsection we compare ArgoCSP-kNN with other CSP portfolios. Since the most important part of our work is the methodology based on very short timeouts (presented in Section 5), and not the portfolio itself, our aim is not to outperform the existing approaches by using ArgoCSP-kNN, but to show that it is comparable with them. The comparison is performed with SUNNY and Proteus, two portfolios that outperformed many other developed approaches.

*Comparison with SUNNY.* ArgoCSP-kNN and SUNNY use different sets of solvers, features and instances, so a direct comparison of the two approaches was not possible without adapting one of them to use the same features, solvers and instances as the other. We applied the SUNNY approach and implemented a SUNNY-like portfolio that uses the same solvers and features<sup>7</sup> as ArgoCSP-kNN and evaluated it on our rich experimental corpus. Apart from different solver scheduling, another difference between SUNNY and ArgoCSP-kNN is that SUNNY uses some fixed number of neighbors  $k$  and some fixed distance measure  $d$ , and our approach determines the optimal values of these metaparameters in the model training phase. The number of neighbors ( $k = 16$ ) that produced the best results in the experiments in the original paper was used. The Euclidean distance measure was used as suggested by the authors. We used 5-fold cross-validation evaluation scheme. The results of comparison between our method and SUNNY are given in Table 2.

Results show that on the used corpus our approach is comparable with SUNNY. Slight improvement may be due to flexibility of our approach with respect to selection of  $k$  and  $d$ , but since the improvement is rather small, we speculate that the situation might even be opposite on different instances, and that no significant difference in performance exists.

---

<sup>7</sup> We also experimented with `mzn2feat` tool for collecting features provided by the authors of SUNNY, but it was very slow on the instances used in our experiments and could not collect features of all instances even after several days.

	Solved instances (out of 8436)
<i>best-fixed</i>	6216
SUNNY	7493
ArgoCSP-kNN	7511
<i>oracle</i>	7631

**Table 2** Results of experimental comparison of ArgoCSP-kNN and SUNNY.

	Solved instances (out of 1493)
<i>best-fixed</i>	984
ArgoCSP-kNN	1413
<i>Proteus</i>	1424
<i>oracle</i>	1493

**Table 3** Results of experimental comparison of ArgoCSP-kNN and Proteus.

*Comparison with Proteus.* As features of used instances and the times for all solvers on all instances used in Proteus paper (Hurley et al. 2014) had been made available online by the authors, we decided to apply the ArgoCSP-kNN approach to that data. All instances in the Proteus corpus satisfy two criteria: they are not trivially solved during 2 seconds of feature computation and they are solved by at least one of the used solvers within the time limit of 1 hour. The set of solvers included 4 CSP solvers and 6 SAT solvers, and three different encodings (direct, support and order) were used when reducing to SAT. For ArgoCSP-kNN, we did not use features of generated SAT instances but only 36 features obtained from the input CSP instances. The results of comparison between ArgoCSP-kNN and Proteus are given in Table 3. The number of solved instances of Proteus is taken from the paper introducing this portfolio. As in the case of SUNNY the difference is very small, so we assume that two systems perform roughly the same.

Note that ArgoCSP-kNN solved respectively 98.4% and 94.6% of possibly solvable instances in experimental results presented in Table 2 and Table 3. Considering the closeness to the oracle, we consider all of the evaluated approaches to be state-of-the-art.


#### 4.4 The effect of the solving timeout on the portfolio effectiveness

The bulk of the time for building the ArgoCSP-kNN portfolio is spent in the preparation phase (instance solving) and this time is directly dependent on the value of the solving timeout used in that phase. Since this is performed only once, usually we are prepared to allow more time for preparation in order to achieve better results in the portfolio exploitation. However, we wanted to examine the effect of the solving timeout on the overall effectiveness of ArgoCSP-kNN. In all cases, the effectiveness is evaluated using 5-fold cross validation over the whole corpus. The preparation is repeated several times with different values of solving timeouts, ranging from 1 to 600 seconds, and to get a fair comparison, during cross-validation the corpus was always partitioned to the same 5 folds. Regardless of the solving timeout used in preparation, the time given to the selected solver during exploitation was 600 seconds.

Table 4 shows the obtained results and they are plotted in Figure 3 (with results for additional values of timeouts). The figure clearly shows that in the beginning there is a

Solving timeout (seconds)	1	5	10	30	60	90	120	300	600
Inst. solved in exploitation	6918	7051	7103	7328	7392	7422	7469	7495	7511
Prep. & train. time (days)	1.4	5.8	10.8	28.9	53.8	77.2	99.7	219.3	393.2
Exploitation time (days)	11.8	11.0	10.5	8.8	8.4	8.2	8.0	8.0	7.9
Total time (days)	13.2	16.8	21.3	37.7	62.2	85.4	107.7	227.2	401.2

**Table 4** Results of experimental evaluation of ArgoCSP-kNN using different solving timeouts (given in seconds). Time spent is given in days.



**Fig. 3** Results of experimental evaluation of ArgoCSP-kNN compared to the *best-fixed* solver and *oracle*; each mark on curve represents one outcome based on solving using the corresponding timeout.

sharp increase in the number of solved instances with the increase of solving timeout, but after certain point the curve stabilizes. E.g., for solving timeouts of 600 and 30 seconds the difference in number of solved instances is not very big (especially when their performance is compared to method *best-fixed*), while the time used for preparation and training reduces from 393.2 to 28.9 days. This demonstrates that it is possible to achieve almost the same results even with much shorter solving timeouts.


## 5 Short Portfolio Training

Users often aim at solving as many instances from a single fixed set of instances, as fast as possible, with the given set of available solvers. A straightforward way would be to choose a single solver that is expected to give the best results (e.g., choose the solver that gave the best results on some solver competition) and apply it on all instances. However, it is not clear how to choose a solver and the chosen solver might not be suited to the specific set of instances. Another approach would be to use a previously trained portfolio. However, the portfolio might have been trained on the set of instances of significantly different properties than the ones that need to be solved, so it might perform poorly.

Encouraged by the results of the previous section that showed that one can significantly reduce the solving timeout without significantly reducing the overall portfolio quality, we propose a new way to efficiently use portfolio approaches in situations when there is a need for solving a previously unseen corpus as fast as possible, using a set of available solvers. The approach relies on short portfolio training on the very same corpus of instances that need to be solved. Due to the nature of the problem (only solving single fixed corpus is

Solving timeout (seconds)	1	5	10	30	60	90	120	300	600
Solved instances	6925	7088	7145	7377	7444	7490	7536	7604	7631
Prep. & training time (days)	1.4	5.8	10.8	28.9	53.8	77.2	99.7	219.3	393.2
Exploitation time (days)	11.8	10.6	10.1	8.2	7.5	7.1	6.8	6.1	0
Total time (days)	13.2	16.4	20.9	37.1	61.3	84.3	106.5	225.4	393.2

**Table 5** Results of evaluation of ArgoCSP-kNN using short training approach for different solving timeouts (given in seconds). The total time used is given in days.



**Fig. 4** Results of evaluation of ArgoCSP-kNN using short training compared to the *best-fixed* solver and *oracle*. Each mark on the curve represents one outcome based on corresponding solving timeout.

important), the generalization of the portfolio to instances out of this corpus is not required, so the overlapping between the training and the test set of instances, which is most common concern in machine learning, becomes irrelevant and will be disregarded in the rest of this section.

The approach is simple and it consists of applying all solvers to all instances in the corpus with a given short solving timeout, performing predictive model training, and then applying the portfolio to the corpus containing only instances not solved during preparation, giving the selected solvers higher solving timeout in exploitation phase.

## 5.1 Evaluation

In the evaluation, we used different solving timeouts in training phase, but due to the nature of our scenario (we cannot afford long solving time), we focused on using very short timeout values. The results are shown in Table 5 and plotted in Figure 4. The total number of solved instances rapidly increases in the region of short solving timeouts and then stabilizes. For the solving timeout of 600 seconds, there is no need to run selected solver on instances not solved in preparation phase with that very same timeout (everything is done in the preparation phase and the number of solved instances is the same as for the *oracle* solver).


The results in tables 4 and 5 and curves in figures 3 and 4 might seem similar. The key difference is that the evaluation in Table 4 and Figure 3 is done by using cross-validation and there is no overlapping between the training and the testing set of instances. Contrary to that, in the evaluation shown in Table 5 and Figure 4 the same corpus is used both for



training and testing (if an instance is solved during training, it is not solved again during testing). Therefore, the short training method constantly achieves better performance, but this is expected as the training and the test sets consist of the same set of instances.

In the practical setting that we described, central quality measures are the number of solved instances and the total time spent for both training and exploitation – one should try to maximize the number of solved instances but keep the total time spent short (e.g., so that both portfolio training and exploitation can be performed on a computer that user has on his disposal). Surprisingly, our experiments show that even for extremely small solving timeouts, the number of solved instances is greater than for the *best-fixed* solver (e.g., for the solving timeout of only 1 second, the number of solved instances increases from 6216 to 6925, while the total time reduces from 16.9 to 13.2 days, where 1.4 days are used for training and 11.8 are used for exploitation). In approximately the same time used by the *best-fixed* solver to solve 6216 instances, our approach solves 7088 instances (for solving timeout of 5 seconds). The relation between solving timeout and total time used is presented in Figure 5). The figure shows that total time used grows linearly with the increase of the timeout. We argue that in practical scenario it makes sense to use solving timeout of up to 30 seconds, depending on the available time, because the number of solved instances does not grow as fast for larger timeout values.

Another important question arises — how much performance do we lose using the proposed approach compared to the situation in which we already have a portfolio built on the instances of similar properties to the instances being solved? A good estimate is readily provided by comparison with the results of evaluation performed in the previous section and presented in Table 4. If ArgoCSP-kNN was trained with the solving timeout of 600 seconds, it would need 393.2 days for training and it would solve 7511 instances. On the other hand, if ArgoCSP-kNN was used with the short training method with the solving timeout of 60 seconds and then with exploitation timeout of 600 seconds per instance, the portfolio would need 53.8 days for training and it would solve 7444 instances (Table 5). The difference of 11 months in training time is obviously very significant and the difference in number of solved instances, maybe not as much, but that is left to the user to decide, based on the purpose and the time available.



**Fig. 5** Results of evaluation of ArgoCSP-kNN using short training and the *best-fixed* solver. Each mark on curve represents one outcome based on corresponding solving timeout.

## 5.2 Comparison of different machine learning techniques

K-nearest neighbors is one of the simplest machine learning techniques. Therefore, it is interesting to perform a comparison to some other machine learning techniques and see if k-nearest neighbors should be substituted by a more promising method in the context of short-training. The way we use k-nearest neighbors algorithm in ArgoCSP-kNN is a bit specific to particular application, and not all machine learning algorithms can be applied in the same way. Therefore, in the comparison we use a bit different, but perfectly natural portfolio design (e.g. SATzilla uses such design (Xu et al. 2008)), and also include ArgoCSP-kNN as it is presented in previous subsection.

The portfolio design used in this subsection is based on prediction of solving time. In the exploitation phase, for each solver, its runtime on input instance is predicted by previously learned model and the solver with the least predicted runtime is invoked to solve the input instance. In the model training phase, the features of all training instances and PAR10 scores of all CSP solvers on all of these instances are used to train the model for each solver. In order to build a model, several combinations of values of metaparameters are evaluated on the training set using 5-fold cross-validation. The best combination is used to build a model on the training set for each solver.

We conducted experiments with regression machine learning techniques to evaluate their efficiency. We experimented with: *k*-NN<sup>8</sup>, linear regression and SVM. Tool RapidMiner<sup>8</sup> was used for the experiments. As RapidMiner performed very poorly when used with SVM (the process just got stuck with many different combinations of metaparameters), we also experimented with tool LIBSVM (Chang and Lin 2011) which showed greater stability. We used two types of kernels with SVM: linear and RBF.

For the efficiency reasons, we used subset of 28 features from 70 features (we eliminated almost half of the features because they had value 0 for all instances and a few features that were in our opinion less important). We used 5-fold cross-validation. The training set and the testing set were normalized before usage. Given a training set, all its instances are solved using each of the included solvers in a given solving timeout, and then the optimal metaparameters are selected in the way already described. The metaparameters for *k*-NN were the number of neighbors *k* (ranging from 1 to 20) and the distance measure *d* (4 different distance measures described by Tomović et al. (Tomović et al. 2006)). The metaparameters for SVM using linear kernel were *C* (powers of 2 from  $2^{-6}$  to  $2^{15}$ ), *v* (ranging from 0.1 to 0.9 with step 0.1). For RBF kernel we used the same combinations of metaparameters as for linear kernel and additionally metaparameter  $\gamma$  (powers of 2 from  $2^{-15}$  to  $2^5$ ). The metaparameters for linear regression were *ridge* (powers of 10 from  $10^{-8}$  to  $10^{-1}$ ), *use\_bias* (true or false), *feature\_selection* (none, M5 prime or greedy) and *eliminate\_colinear\_features* (true or false). For each method, all combinations of metaparameters are tried on the training set and the ones generating the best score are declared optimal.

Table 6 shows the obtained results. The *k*-NN method solves significantly more instances than other used approaches, and ArgoCSP-kNN solves even more (its performance with the reduced set of features is almost the same as with the full set, already presented in Table 5). As idea is to compare different machine learning techniques using short training, we had to limit the time used for prediction. For both kernels in libsvm tests, months would be needed to finish these experiments, so we used the reduced sets of metaparameter values giving the best results in preliminary experiments. Even with these reduced set of parameters libsvm

---

<sup>8</sup> <https://rapidminer.com/>

tests with RBF kernel needed more than a month to finish. We therefore conclude that due to time consumption, SVM models are not suitable for short training approach.

	1	5	10	30	60
Linear regression (RM)	5947	5950	6073	6291	6459
Linear kernel (libsvm)	6216	6753	6753	6762	6695
RBF kernel (libsvm)	6235	6663	6518	6716	6748
$k$ -NN (RM)	6052	6644	6804	7201	7309
$k$ -NN (ArgoCSP-kNN)	6927	7095	7147	7375	7458

**Table 6** Results of experimental evaluation using different machine learning techniques and different solving timeouts.

As in subsection 4.4, we estimate performance of portfolio previously trained on similar (but not the same) set of instances, using 5-fold cross-validation on the corpus with different machine learning techniques and solving timeout of 600 seconds. Results are given in Table 7. In this case, no reduction of metaparameters in case of SVM was performed, but still, both versions of  $k$ -NN perform better than other approaches.


	Solved instances
<i>best-fixed</i>	6216
Linear regression (RM)	6515
Linear kernel (libsvm)	6715
Rbf kernel (libsvm)	7323
$k$ -NN (RM)	7432
$k$ -NN (ArgoCSP-kNN)	7511
<i>oracle</i>	7631

**Table 7** Results of evaluation using 5-fold cross-validation and different machine learning techniques.

### 5.3 Evaluation on SAT Instances

Since our short training methodology gave very good results with CSP portfolio inspired by ArgoSmArT-kNN, it is interesting to show that our short training methodology is applicable to SAT corpora, as well. Therefore we run experiments on SAT instances using original ArgoSmArT-kNN portfolio.

We compare ArgoSmArT-kNN with and without short training methodology on instances originally used in ArgoSmArT-kNN paper (Nikolić et al. 2011) – instances from SAT Competitions (2002-2007) and SAT Races (2000-2008). In short training approach, portfolio is both trained and run on those instances. In original approach without short training, the portfolio is evaluated using 5-fold cross-validation. The results of the comparison are given in Figure 6. Since in the original approach there is no sequence of timeouts, it is represented as a line parallel to horizontal axis. The curve for short training approach looks very similar to the one presented in previous subsection. Very good results are obtained for very small timeouts. Therefore, our methodology can be efficiently applied not only to CSP instances, but also to SAT instances.



**Fig. 6** Results of evaluation of ArgoSmArT-kNN portfolio on SAT instances using short and standard training, *best-fixed* solver and *oracle*. Each mark on curve represents an outcome based on corresponding solving timeout.

## 6 Conclusions and Further Work

In this paper, we have presented a simple  $k$ -NN-based portfolio for CSP and shown that it achieves state-of-the-art performance by comparison with other methods. We have assessed the effects of the training phase duration to overall portfolio quality by using different solving timeout values. Our approach significantly improves over each single constituent solver and gives good generalization results even when used with a very short preparation phase (so that the training does not require an advanced cluster computer, but can be done on a single PC).

Often, solving a single fixed set of instances as fast as possible with the given set of solvers is the only practical concern. We have addressed this problem and proposed a portfolio-based approach that uses short portfolio training on the set of instances to be solved and then selects a suitable solver for each of those instance based on the learned predictive model. Experimental results indicate that this approach significantly improves over the best fixed solver, even for very short timeouts. Also, the number of solved instances first quickly increases when increasing the solving timeout value, but then saturates, and from some point, increasing the solving timeout yields small number of additional solved instances. Therefore, one can choose a solving timeout depending on the time he has available and still significantly improve the results over any single solver.

We also performed experiments with SAT instances and our methodology also shows good results on this kind of problems.

A possible direction in improving short training approach could be to automatically select the solving timeout for the set of instances to be solved, based on the total time available. This would make the approach more autonomous and the one would not need to think about the solving timeout that is going to be used, but to simply determine the total time available. Also, it would be good to consider configuration options of a diverse set of solvers and while choosing a solver also to choose its suitable configuration. For example, systems *meSAT*, *Azucar*, and *BEE* (Metodi and Codish 2012) offer choosing between different

encodings, while Minisat++ (Sörensson and Eén 2009) offers 3 different options when encoding PB constraints (adders, BDDs, Sorters). One more direction of future work, would be to test our approach on other types of problems (e.g. Answer Set Programming, Constraint Optimization Problems).

**Acknowledgements** This work was partially supported by the Serbian Ministry of Science grant 174021.

## References

- Amadini R, Gabbrielli M, Mauro J (2013) An empirical evaluation of portfolios approaches for solving csps. In: CPAIOR, pp 316–324
- Amadini R, Gabbrielli M, Mauro J (2014) SUNNY: a lazy portfolio approach for constraint solving. TPLP 14(4-5):509–524
- Apt KR (2003) Principles of constraint programming. Cambridge University Press
- Beldiceanu N, Carlsson M, Rampon JX (2005) Global constraint catalog. Tech. rep., SICS
- Biere A, Heule M, van Maaren H, Walsh T (eds) (2009) Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications, vol 185, IOS Press
- Bishop CM, et al. (2006) Pattern recognition and machine learning, vol 4. Springer New York
- Bofill M, Suy J, Villaret M (2010) A system for solving constraint satisfaction problems with smt. In: SAT, pp 300–305
- Chang C, Lin C (2011) LIBSVM: A library for support vector machines. ACM TIST 2(3):27
- Cook SA (1971) The complexity of theorem-proving procedures. In: STOC, pp 151–158
- Dutertre B, de Moura L (2006) The Yices SMT solver. Tech. rep., SRI International
- Eén N, Sörensson N (2003) An extensible sat-solver. In: SAT, pp 502–518
- Gent IP (2002) Arc consistency in sat. In: ECAI, pp 121–125
- Hebrard E (2008) Mistral, a constraint satisfaction library. Proceedings of the 3rd International CSP Solver Competition pp 31–39
- Hurley B, Kotthoff L, Malitsky Y, O’Sullivan B (2014) Proteus: A hierarchical portfolio of solvers and transformations. In: CPAIOR, pp 301–317
- Hutter F, Hoos HH, Leyton-Brown K, Stützle T (2009) Paramils: An automatic algorithm configuration framework. J Artif Intell Res (JAIR) 36:267–306
- Kadioglu S, Malitsky Y, Sabharwal A, Samulowitz H, Sellmann M (2011) Algorithm selection and scheduling. In: CP, pp 454–469
- Kiziltan Z, Mandrioli L, O’Sullivan B, Mauro J (2011) A classification-based approach to manage a solver portfolio for csps. In: AICS-2011
- Malitsky Y, Sellmann M (2012) Instance-specific algorithm configuration as a method for non-model-based portfolio generation. In: CPAIOR, pp 244–259
- Malitsky Y, Sabharwal A, Samulowitz H, Sellmann M (2011) Non-model-based algorithm portfolios for sat. In: SAT, pp 369–370
- Malitsky Y, Sabharwal A, Samulowitz H, Sellmann M (2012) Parallel SAT solver selection and scheduling. In: CP, pp 512–526
- Malitsky Y, Sabharwal A, Samulowitz H, Sellmann M (2013) Algorithm portfolios based on cost-sensitive hierarchical clustering. In: IJCAI 2013, AAAI Press, IJCAI ’13, pp 608–614
- Merchez S, Lecoutre C, Boussemart F (2001) Abscon: A prototype to solve csps with abstraction. In: CP, pp 730–744

- Metodi A, Codish M (2012) Compiling finite domain constraints to sat with bee. *TPLP* 12(4-5):465–483
- de Moura LM, Bjørner N (2008) Z3: An efficient smt solver. In: *TACAS*, pp 337–340
- Nethercote N, Stuckey PJ, Becket R, Brand S, Duck GJ, Tack G (2007) Minizinc: Towards a standard cp modelling language. In: *CP*, pp 529–543
- Nikolić M, Marić F, Janičić P (2009) Instance-based selection of policies for sat solvers. In: *SAT*, pp 326–340
- Nikolić M, Marić F, Janičić P (2011) Simple algorithm portfolio for sat. *Artificial Intelligence Review* pp 1–9
- Ohrimenko O, Stuckey PJ, Codish M (2009) Propagation via lazy clause generation. *Constraints* 14(3):357–391
- O’Mahony E, Hebrard E, Holland A, Nugent C, O’Sullivan B (2008) Using case-based reasoning in an algorithm portfolio for constraint solving. In: *AICS-2008*
- Régin JC (2004) Global constraints and filtering algorithms. In: *Constraint and Integer Programming*, Springer, pp 89–135
- Rossi F, van Beek P, Walsh T (eds) (2006) *Handbook of Constraint Programming*. Elsevier
- Roussel O, Lecoutre C (2009) Xml representation of constraint networks: Format xcsp 2.1. *CoRR* abs/0902.2362
- Schulte C, Lagerkvist M, Tack G (2006) Gecode. Software download and online material at the website: <http://www.gecode.org>
- Sörensson N, Eén N (2009) Minisat 2.1 and minisat++ 1.0sat race 2008 editions. *SAT* p 31
- Stojadinović M, Marić F (2014) mesat: multiple encodings of CSP to SAT. *Constraints* 19(4):380–403
- Tamura N, Banbara M (2008) Sugar: A csp to sat translator based on order encoding. In: *Proceedings of the third constraint solver competition*, pp 65–69
- Tamura N, Taga A, Kitagawa S, Banbara M (2009) Compiling finite linear csp into sat. *Constraints* 14(2):254–272
- Tanjo T, Tamura N, Banbara M (2011) A compact and efficient sat-encoding of finite domain csp. In: *SAT*, pp 375–376
- Tanjo T, Tamura N, Banbara M (2012) Azucar: A sat-based csp solver using compact order encoding - (tool presentation). In: *SAT*, pp 456–462
- Tomović A, Janičić P, Kešelj V (2006) n-gram-based classification and unsupervised hierarchical clustering of genome sequences. *Computer Methods and Programs in Biomedicine* 81(2):137–153
- Walsh T (2000) Sat v csp. In: *CP*, pp 441–456
- Xu L, Hutter F, Hoos HH, Leyton-Brown K (2008) Satzilla: Portfolio-based algorithm selection for sat. *J Artif Intell Res (JAIR)* 32:565–606